Inferring Object Properties from Incidental Contact with a Tactile-Sensing Forearm

Tapomayukh Bhattacharjee*, James M. Rehg, and Charles C. Kemp

arXiv:1409.4972v1 [cs.RO] 17 Sep 2014

***Abstract*—Whole-arm tactile sensing enables a robot to sense properties of contact across its entire arm. By using this large sensing area, a robot has the potential to acquire useful information from incidental contact that occurs while performing a task. Within this paper, we demonstrate that data-driven methods can be used to infer mechanical properties of objects from incidental contact with a robot's forearm. We collected data from a tactile-sensing forearm as it made contact with various objects during a simple reaching motion. We then used hidden Markov models (HMMs) to infer two object properties (rigid vs. soft and fixed vs. movable) based on low-dimensional features of time-varying tactile sensor data (maximum force, contact area, and contact motion). A key issue is the extent to which data-driven methods can generalize to robot actions that differ from those used during training. To investigate this issue, we developed an idealized mechanical model of a robot with a compliant joint making contact with an object. This model provides intuition for the classification problem. We also conducted tests in which we varied the robot arm's velocity and joint stiffness. We found that, in contrast to our previous methods [1], multivariate HMMs achieved high cross-validation accuracy and successfully generalized what they had learned to new robot motions with distinct velocities and joint stiffnesses.**

***Index Terms*—Haptics, Tactile Sensing, Hidden Markov Models, Classification.**

I. INTRODUCTION

Manipulation in unstructured environments with high clutter is difficult due to a variety of factors, including uncertainty about the state of the world, a lack of non-contact trajectories, and reduced visibility for line-of-sight sensors [2]. Tactile sensing is well-matched to these challenges, since it benefits from contact and uses sensors that move with the manipulator into the clutter. When contact occurs with the tactile sensors the robot often has an opportunity to acquire useful information. By fully covering the robot's manipulator with tactile sensors, the robot is likely to have more opportunities to acquire useful information through contact. However, with a typical serial manipulator, a robot cannot independently control the pose of each of the sensors and contact can be unanticipated.

Within this paper, we address the problem of tactile perception based on incidental contact [1], [3], [4] with a tactile-sensing forearm. By *incidental contact*, we mean contact that is not central to the robot's current actions and may occur unexpectedly or unintentionally.[1] As such, incidental contact will typically not involve active exploration and interrogation of the contact, since the robot will be directing its resources elsewhere.

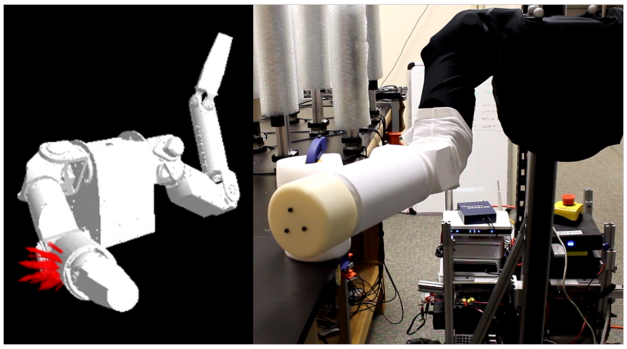

Fig. 1: *Force data from forearm skin sensor mounted on Cody while the forearm is pushing against a fixed cylindrical object made of polystyrene foam. The red arrows show the force acting on the skin.*

We have previously provided evidence for the value of allowing contact, even unintended contact, to occur with the robot's manipulator [2], [5]. We have also found that contact with a robot's forearm can be relatively frequent when reaching in clutter or assisting a person in a wheelchair [5], [6]. In this paper, we investigate the potential for data-driven methods to enable a compliant robot with a tactile-sensing forearm to infer two mechanical properties of an object (rigid vs. soft and fixed vs. movable) based on contact during a simple reaching motion. In Section III, we present our method, which uses multivariate HMMs to infer object properties. HMMs have a long history of success for classifying time series such as human speech [7], and we have previously shown that HMMs can rapidly recognize the type of an object in a known environment based on tactile sensing during incidental contact [3]. In Section IV-A, we show that our method had high cross-validation accuracy when evaluated with tactile data from a robot making contact with 18 different objects.

Inferring mechanical properties of objects from incidental contact could be beneficial in a number of ways. For example, we have shown that haptically recognizing leaves vs. trunk while reaching into artificial foliage can be used by a robot to haptically map the environment and plan paths to goals [4]. Rather than recognizing a particular object type, detecting an object's properties could be advantageous in novel environ-

T. Bhattacharjee and C. C. Kemp are with the Healthcare Robotics Lab, Center for Robotics and Intelligent Machines, Georgia Institute of Technology,

J. M. Rehg is with the Computational Perception Lab, Center for Robotics and Intelligent Machines, Georgia Institute of Technology,

*T. Bhattacharjee is the corresponding author {tapomayukh@gatech.edu}.

[1]This description supersedes our previous descriptions from [3] and [4].

Fig. 2: *Set of objects for experiments with stereotyped motion.*

ments. Detecting that an object is movable could be used by the robot to make better decisions, such as moving the object in order to access a new location or avoiding the object so as not to alter the environment. Likewise, detecting that an object is rigid or soft has implications for the robot's ability to compress the object and the consequences of collisions with the object. Tactile sensing is well-suited to the perception of mechanical properties of objects, since it can directly sense forces and monitor the locations of contact over time.

The signals produced from a robot's tactile sensors depend on the mechanics of both the object and the robot. As such, a key issue for data-driven approaches to tactile perception of incidental contact is the extent to which perceptual classifiers can perform well when a robot's actions differ from those used during training. To gain intuition for these issues, we present an idealized mechanical model of a robot with a compliant joint making contact with an object in Section V. Then, in Section VI, we present tests with a real robot for which we varied the robot arm's velocity and joint stiffness to values distinct from those used during training. As we discuss in Section VI-A, our HMM-based method performed well in these tests, but a previous method of ours from [1] performed poorly. This older method used the k-nearest neighbor algorithm (k-NN) to classify vectorized and dimensionality-reduced time series of features (maximum force, contact area, and contact motion). Among other challenges, scaling a signal in time can result in large changes in this vector representation, resulting in large distances between vectors representing interactions with similar objects. In Section IV, we present additional results from tests in which we varied the perceptual features and parameters used by our method.

II. RELATED WORK

Object categorization is a well-studied task. We focus on the haptic sensing modality in this work. Although there have been multiple studies on haptic-based compliance discrimination, most have used specific exploratory behaviors using end effectors to extract information from the environment. Studies of discrimination tasks using information from *incidental contact through large-area tactile sensing* are lacking. Furthermore, work focusing on predicting object mobility using haptics is rare. In the following subsections, we review the existing literature that addresses the object categorization task using haptics.

A. Material Property based Classification

Previous work on material property classification is perhaps the most closely related work to ours. Although we do not explicitly model material properties, the features that we extract from the interactions between the robot arm and environmental objects are a direct consequence of these material properties that affect the interaction dynamics. Drimus *et. al.* [8] classified rigid and deformable objects based on haptic feedback from a novel tactile sensor comprised of a flexible, piezo-resistive rubber. They represented tactile information from a palpation procedure as a time series of features and used a k-nearest-neighbor (k-NN) classifier to categorize the objects [8]. Our classification scheme considers both compliance and mobility characteristics and uses information from incidental contact sensed with large-area tactile sensors. In addition, the features extracted in our method correspond to physical quantities that more clearly relate the underlying mechanics.

Jain *et. al.* used data-driven object centric models to haptically recognize specific doors as well as classes of doors (refrigerator vs. kitchen cabinet) [9]. Chu and McMahon *et. al.* [10] as well as McMahon and Chu *et. al.* [11] present research that uses HMMs to automatically assign adjectives to haptic signals collected while a robot interacts with an object using exploratory behaviors. In contrast to our work, their research focuses on classifying long time series (≈80 s) based on deliberate and focused probing with sophisticated robotic fingers from *Syntouch BioTacs*.

Sukhoy *et. al.* investigated the use of a vibro-tactile sensor and a support vector machine (SVM) classifier for surface texture recognition [12]. Kim and Kesavadas presented a methodology for estimating the material properties of objects using an active tapping procedure [13]. Takamuku *et. al.* estimated the material properties of objects through tapping and squeezing behaviors [14]. Hosoda and Iwase [15] used a bionic hand and used its compliance to grip an object to obtain haptic data. They used a recurrent neural network to classify objects based on haptic cues learned from dynamic interactions [15]. Nizar *et. al.* [16] classified the material type and surface properties by developing a sensor that uses a lightweight plunger probe to detect surface properties. They also used an optical mouse sensor to obtain surface images and used a radial basis function neural network for classification.

In summary, although there have been many studies on material-property-based classification, most of them have focused on specific exploratory behaviors using the robot's end-effector.

B. Shape-based Classification

Many researchers have used tactile images from touch sensors and employed different algorithms for object identification. Schneider *et. al.* [17] applied a "bag-of-words" approach and unsupervised clustering techniques to categorize

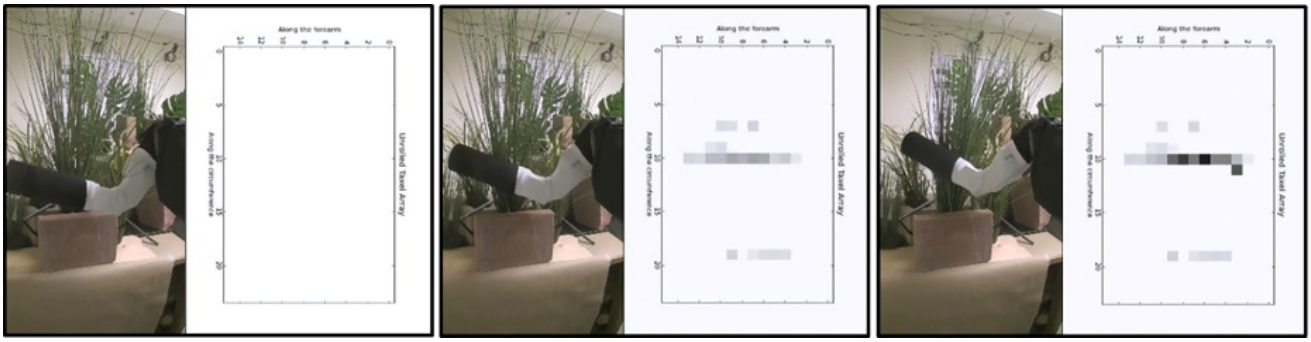

Fig. 3: *Sequence of images that illustrates the data collection for our experiments on inferring mechanical properties of objects (foliage). Each image shows a picture of the robot Cody and a visualization of the data from the forearm skin sensor as a 24X16 image (dark pixels correspond to larger forces). The leftmost picture shows a non-contact situation, the middle picture corresponds to the situation just after the onset of contact, and the rightmost picture shows the situation when the robot has pushed the foliage to the maximum extent.*

objects. Pezzementi *et. al.* identified the principal components of identified features, then clustered them, and constructed per-class histograms as a class characteristic [18]. Gorges *et. al.* [19] introduced passive joints in the hand for better adaptibility to different object shapes and used a Bayes classifier to classify the objects.

Some other researchers analyze the deformation behavior to classify objects. They use vision and haptic sensors [20] or finite element models [21], [22] or volumetric models such as superquadrics [23], polyhedral models [24] and wrapping polyhedron [25] to match recovered model features after grasping with multi-fingered hands.

To summarize, shape-based classification schemes have used tactile images or deformation behaviors to classify objects after exploring or grasping them using multi-fingered robot hands.

C. Functional Property based Classification

This group of studies focused on functional property/behavior-based classification methods. Sinapov *et. al.* used the acoustic properties of objects during specific interaction schemes and the behavioral interactions performed with them, such as grasping, shaking, dropping, pushing, and tapping, to classify 36 different household objects [26]. Berquist *et. al.* monitored the changes in the joint torques of a robot while it performed five exploratory procedures - lifting, shaking, crushing, dropping, and pushing - on several objects and demonstrated that the robot can learn to recognize objects based solely on the joint-torque information [27]. Griffith *et. al.* used multiple exploratory behaviors and employed clustering techniques to categorize containers and non-containers. After extracting visual and acoustic features from the interactions with objects, they employed unsupervised clustering techniques to form several categories [28]. Sinapov *et. al.* combined proprioceptive and auditory feedback and used a behavior-grounded relational classification model to recognize categories of household objects [29].

III. METHODS

We used supervised machine learning methods to process and extract information using data from a skin sensor covering the forearm of a humanoid robot named "Cody". Our goal was to classify an object as being in one of four categories: *Rigid-Fixed* (RF), 2) *Rigid-Movable* (RM), 3) *Soft-Fixed* (SF), and 4) *Soft-Movable* (SM). In Sec. IV-B, we show the effect of the spatial resolution of the taxels on the classification accuracy. Section IV-C highlights the importance of the number of hidden states for classification purposes.

A. Algorithm

Our method of classification involves some preprocessing and feature selection of the raw data as described in Section III-A1, and then uses HMMs to model the time-series of the selected features as explained in Section III-A2.

1) Preprocessing and Feature Selection: The experimental setup and data collection procedure are detailed in Section III-B. After recording the time-series data using the forearm taxel array (details in Section III-B), we truncated them to begin at the estimated onset of contact (whenever the force exceeds a threshold) between the robot and object. We then represented the data at every time step as a gray-scale image, as shown in Fig. 3. We converted this image to a binary image, representing the taxels in contact by applying a threshold to each taxel. This threshold was not the same for all objects. We selected a higher threshold for some of the more rigid or coarser objects to account for an extra covering that we put over the otherwise bare tactile sensors to protect them from damage. Then, we computed the connected components to segment the contact regions. For the connected component with the largest area, we computed three features. The flowchart shown in Fig. 4 illustrates the preprocessing we performed.

Fig. 5 shows an example of how the three features vary over time during incidental contact. The first feature is the maximum force (F_{max}) measured by a taxel in the contact region at each time step. This is analogous to measuring the highest pressure. Initial experiments show that maximum force performs better than total force or mean force. This is probably

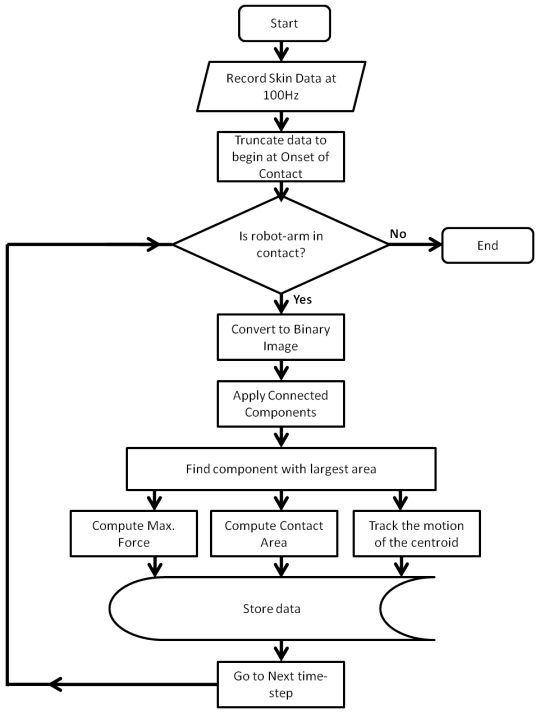

Fig. 4: *Flowchart of the preprocessing steps of the raw skin data.*

because total force would not be able to distinguish cases of small forces over large contact area with cases of large forces over small contact area. Mean force in a contact area would also not be able to distinguish large forces if the surrounding area has low forces. Second, we estimated the area of contact (a) between the arm and object (contact region) as the number of taxels in the connected component. Third, we estimated the distance that the centroid of the connected component traveled in the world frame from its position at the onset of contact (d). We assumed that the robot's torso did not move throughout the trials and used the forward kinematics from the robot's torso to the center of the contact location on the robot's forearm to estimate the 3D positions and distance. We expected these three features to be informative about the object's softness and mobility. For example, with increasing force applied to a soft, fixed object, we would expect the contact area to increase. Likewise, we would expect the 3D position of the contact area to travel when encountering movable and soft objects. When making contact with a rigid and fixed object, we would expect the maximum force to increase. We obtained the values of maximum force, the number of taxels in the contact region, and the contact motion for each trial during the first 1.2 s time window after the onset of contact.

Based on our previous results in [1], we expected that the maximum force (F_{max}) and contact area (a) features would be important, and thus, we used univariate HMMs (details in Section III-A2) to model the temporal trends of these feature vectors for categorization. Figure 6 depicts the complete experimental protocol.

2) Hidden Markov Model Implementation: Hidden Markov model is a state-based data modeling tool that assumes the states are hidden and the current state is dependent only on the previous state. The hidden states are inferred using observations. The components of an HMM include (1) N, the number of states in the model; (2) A = $\{a_{ij}\}$, the state transition probabilities; (3) B = $\{b_j\,(k)\}$, the observation probabilities; and (4) P = $\{\pi_i\}$, the initial state probabilities [7]. The HMM model (λ) is represented as shown in eq. (1). The state transition probabilities are shown in eq. (2) and the initial state probabilities are represented in eq. (3).

$$\lambda = (A, B, \pi) \tag{1}$$

$$A = \{a_{ij}\} = \{P\,(x_t = j | x_{t-1} = i)\} \tag{2}$$

$$\pi = \{\pi_i\} = \{P\,(x_0 = i)\} \tag{3}$$

For a discrete HMM, the observation probabilities are shown in eq. (4).

$$B = \{b_j\,(k)\} = \{P\,(o_t = k | x_t = j)\} \tag{4}$$

However, for a continuous HMM with gaussian output probability density functions, the observation probabilities are given by eq. (5).

$$B = \{b_j\,(k)\} = \{P\,(o_t | x_t = j)\} = \{f\,(o_t, (\mu_j, \sigma_j))\} \tag{5}$$

We implemented univariate continuous HMMs to model the temporal trends of maximum force and contact area for different categories of objects. Our objective was to use these models to determine whether they can characterize the differences between each of the four object categories and classify them accordingly using information from incidental contact. We modeled each of these four object categories- *Rigid-Fixed*, *Rigid-Movable*, *Soft-Fixed*, and *Soft-Movable*- using eight HMMs (4 categories X 2 features = 8 HMMs). Each of these HMMs is a left-right HMM with 10 hidden states. In a left-right HMM, the transitions from one state to another can only go forward in one direction, which is analogous to the progression of our data in feature space. For each of these categories, one of the HMMs models the maximum force, and the other models the contact area. We used a 1D HMM for each feature (maximum force and contact area) which helped us analyze the effects of each feature for classification purposes. We time-normalized the data and sampled them into N equal parts (corresponding to the number of states N). For each part, we expressed all the data for a specific feature and a particular category using a Gaussian distribution with a mean, μ and a standard deviation, σ. Therefore, we have N Gaussian distributions from the observed data for each HMM model. These represent our emissions.

We obtained the transition probabilities for each of these HMMs by training them on the experimental data. We used the GHMM toolkit [30] to model the HMMs and implement them in Python. We trained the models with the standard Baum-Welch algorithm, which uses expectation maximization. For

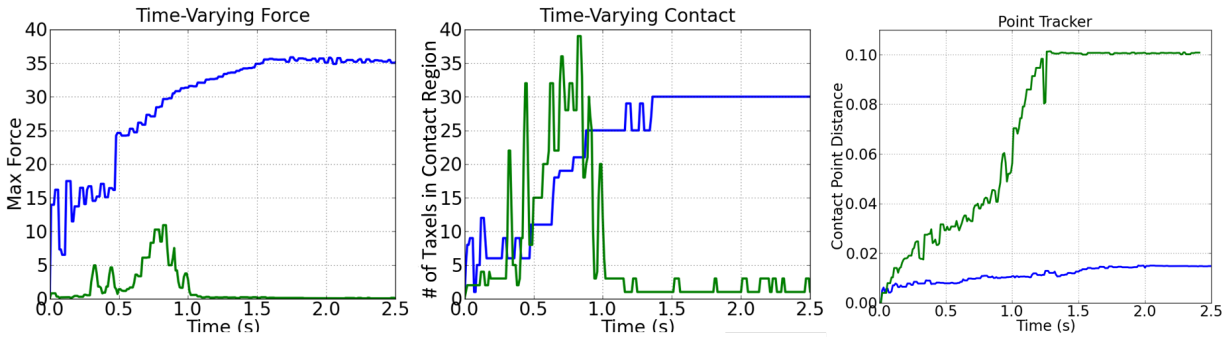

Fig. 5: *Example of the three features that we computed from the data from the forearm skin sensor and used to classify object properties. The leftmost picture shows the maximum force over time (in Newtons), the middle picture shows the contact area over time, and the rightmost picture shows the contact motion (in meters). The green lines are the features for a movable object, and the blue lines are the features for a fixed object.*

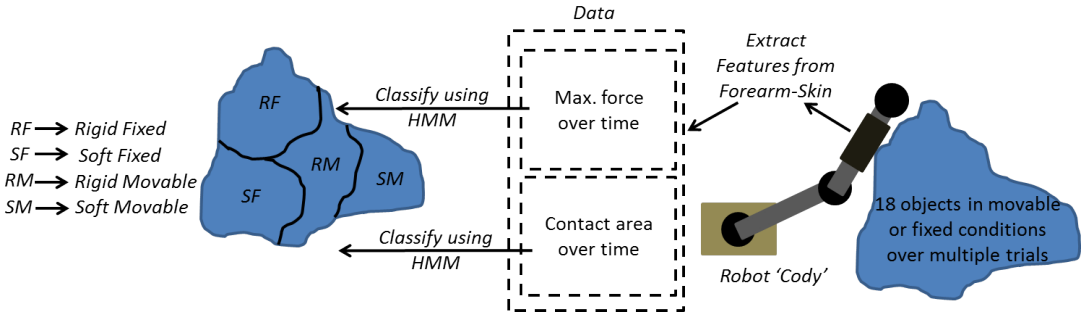

Fig. 6: *Schematic representation of the experimental protocol.*

testing, we ran the Viterbi algorithm which estimates the most probable state sequence for given test data and then finds the model that represents that state sequence with the highest probability. These are standard methods for modeling sequential data, see [7] for details. We ran the Viterbi algorithm on the given test data for all the trained 8 HMM models and identified the model that returned the highest probability. Figure 7 shows the schematic of one HMM model. Our entire algorithm is shown using the pseudo-code given in Algorithm 1.

B. Experimental Procedure : Experiments with Stereotyped Motion

Section III-B1 describes the experimental setup while Section III-B2 explains our data collection method.

1) Experimental Setup: Cody, as shown in Fig. 1, is a statically stable mobile manipulator weighing approximately 160 kg. The components of the robot are: Meka A1 arms, a Segway omni-directional base and a Festo 1-DOF (degree of freedom). linear actuator. The two seven-DOF anthropomorphic arms contain series elastic actuators. When we control these arms, each joint simulates a low-stiffness, visco-elastic, torsional spring. We control the robot's arms by changing the equilibrium angles of these simulated springs over time.

Cody has a force-sensitive skin across its entire forearm. Meka Robotics and the Georgia Tech Healthcare Robotics Lab developed the forearm tactile skin sensor, which is based on Stanford's capacitive sensing technology, as described by Ulmen *et. al.* [31]. The skin consists of a capacitive pressure-sensor array. We refer to the elements of this array as taxels (tactile pixels). There are 384 taxels on the entire skin, and these are distributed in a 24X16 array, with each taxel being 9 mm X 9 mm in size. The array of taxels reports the measured force applied to each taxel at 100 Hz.

2) Data Collection: For our experiments, we used a set of 18 objects, as shown in Fig. 2. We selected large objects that have mostly uniform material properties and vary widely in their mass, friction, and compliance. We actuated the robot's elbow joint by commanding a goal point in the joint space. The robot arm tries to reach the goal using a joint PD controller. The final goal point in joint space is selected such that the equivalent point in the Cartesian space is inside the object. Thus, the robot will come in contact with the object before reaching the goal. When the robot incidentally comes in contact with the object, it pushes against it and tries to reach the goal as shown in Fig. 3. For each object, we collected haptic data by commanding the same goal position for the arm and recording the sensor readings from the taxels of the forearm skin at approximately 100 Hz. We labeled each of these objects as either soft or rigid. We considered pillow-like materials, such as foam, bubble wrap, and vegetation, to be soft and all other objects to be rigid. For objects that could be pushed aside by the robot's motion, we fixed them with a clamp or a heavy weight so that we could have both movable and fixed conditions. We collected data from five trials for each of the 18 different objects, 10 of them in both fixed and

Algorithm 1 Classification Algorithm

$F_{list}, a_{list}, d_{list} \leftarrow [\,]$
$i \leftarrow 0$
$p_0 \leftarrow position\ of\ the\ arm\ at\ onset\ of\ contact$
while robot in contact **do**
  $F_{max}, a, d \leftarrow 0$
  $C_i \leftarrow identify\ connected\ components$
  $C_L \leftarrow largest\ C_i$
  for all taxels in C_L **do**
    $F_{max} \leftarrow \max(F_{max}, F)$
    $a \leftarrow a + 1$
    $p \leftarrow current\ position\ of\ C_L$
    $d \leftarrow p - p_0$
  end for
  $F_{list}[i] \leftarrow F_{max}$
  $a_{list}[i] \leftarrow a$
  $d_{list}[i] \leftarrow d$
  $i \leftarrow i + 1$
end while
for each category in $[RF, RM, SF, SM]$ **do**
  for each feature in $[F_{list}, a_{list}]$ **do**
    $Train\ \lambda_j\ using\ Baum - Welch\ Algorithm$
  end for
end for
$P_{max} \leftarrow 0$
for new testing time-series data **do**
  $P_i \leftarrow Probability\ from\ Viterbi\ Algorithm$
  $P_{max} \leftarrow \max(P_{max}, P_i)$
end for
for each feature in $[F_{list}, a_{list}]$ **do**
  $\lambda_{max}\ which\ returns\ P_{max}$
  $Test\ data\ belongs\ to\ category\ of\ \lambda_{max}$
end for

movable conditions, four of them in only fixed conditions, and the remaining four in movable conditions. Please note that some of the objects (eg. a heavy iron bucket) could not be moved by the robot's motion even though they were not fixed. Fig. 3 shows three images from one trial of the robot making contact with a plant. It also shows the data from the forearm sensor visualized as an image.

IV. RESULTS AND DISCUSSION

A. Classification Results

We used HMMs to test the classification accuracy for four different classification problems using information obtained from incidental contact. In each case, we used 10 hidden states and expressed the observations as Gaussian distributions in a particular region. We used 10 hidden states and applied five-fold cross-validation to assess the performance.

Fig. 8 presents the confusion matrix for classification into four categories: 1) *Rigid-Fixed*, 2) *Rigid-Movable*, 3) *Soft-Fixed*, and 4) *Soft-Movable*. The classification accuracy was 83.57% with 10 hidden states and using the maximum force as the observed feature. Many of the classification errors were between the *Rigid-Movable* and *Soft-Fixed* classes. Fig.

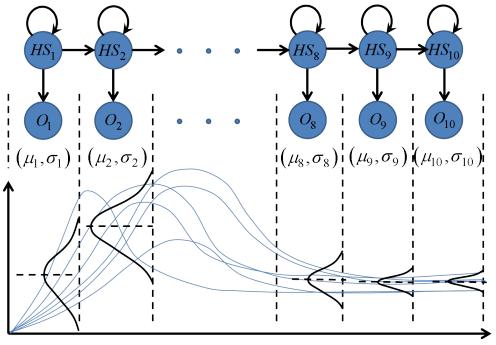

Fig. 7: *Schematic of a HMM with 10 states. The observations are modeled as Gaussian for all of the trials in a particular category.*

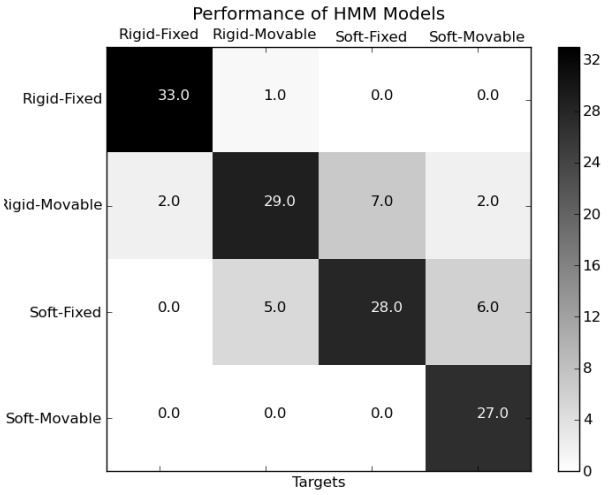

Fig. 8: *Classification into four categories using univariate HMMs with maximum force as observations for experiments with the robot 'Cody'. In these experiments, the robot moved with stereotyped motion.*

9 shows the confusion matrix for the classification with the same number of hidden staes but with contact area as the observed feature. In this case, most errors were between the *Rigid-Movable* and *Rigid-Fixed* classes as well as between the *Soft-Movable* and *Soft-Fixed* classes. The results show that contact area and contact force result in similar performance. The primary source of confusion arises between the fixed and movable categories, suggesting that the contact motion information might be valuable resulting in higher overall performance. *Thus, we decided to implement multivariate Gaussian HMMs with both contact force and contact motion as the observations as these features are both useful.* The detailed implementation and results are provided in Section IV-D.

B. Effect of Taxel Resolution

We performed the four category classification experiment for different spatial resolutions of the taxels. We reduced the taxel resolution by grouping four adjacent (2X2) taxels together into one taxel thereby reducing the taxel numbers

TABLE I: Effect of Taxel Resolution.

$taxels/cm^2$	Classification Accuracy with Force	Classification Accuracy with Area
0.005	78.57%	25%
0.019	80%	40.71%
0.077	73.57%	60.71%
0.309	78.57%	67.86%
1.235	83.57%	84.99%

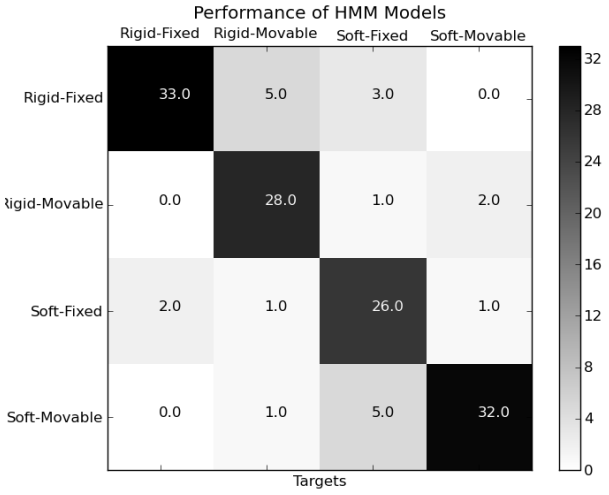

Fig. 9: *Classification into four categories using univariate HMMs with contact area as observations for experiments with the robot 'Cody'. In these experiments, the robot moved with stereotyped motion.*

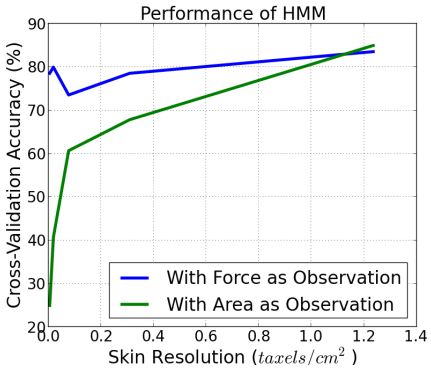

Fig. 10: *As skin taxel resolution increases, the performance increases. The results are from experiments in which the robot moved with stereotyped motion.*

from 384 to 96, 24, and 6 in successive steps. For each of these steps, we added the forces of each of the four taxels to obtain the resultant force for the new taxel. For the last step, we grouped all the 6 taxels into 1 taxel. The resolution is expressed in $taxels/cm^2$ such that for 384 taxels, each taxel was of 9 mm x 9mm size and the size of the taxels increased according to the factor of reduced resolution.

Table I shows the 5-fold cross-validation accuracy for each taxel resolution using maximum force (F_{max}) and contact area (a) as the features. Fig. 10 presents the total and individual category classification accuracies that we obtained for the different resolutions. Compared to 0.005 $taxels/cm^2$, 1.235 $taxels/cm^2$ resolution improved the classification accuracy by 5% using force and by 60% using area. Interestingly, the maximum force feature (F_{max}) for reduced resolution is related to total force because the forces are summed up for neighboring taxels.

C. Effect of States

Finally, we analyzed the effect of the number of hidden states on the performance of the classification algorithm. We ran our algorithm with 10, 15, 20, 30, 50, and 100 states with maximum force (F_{max}) as observations. Fig. 12 shows the results. Table II shows the quantitative results of the four-category classification scheme.

The results in Table II indicate that as the number of states increases, the cross-validation accuracy increases up

TABLE II: Effect of the Number of States.

No. of States	Classification Accuracy with Force as Observations
10 States	83.57%
15 States	89.29%
20 States	91.43%
30 States	91.42%
50 States	89.28%
100 States	73.57%

to 20-30 states, after which it decreases. However, the rate of increase is much higher from 10 to 15 states compared with 15 to 20 and 20 to 30. Increasing the number of states beyond that level might have resulted in over-fitting and decreased performance. The performance increase from 10 to 20-30 states is approximately 8% but comes with the burden of greater computational expense. A task-based, application-specific choice of the number of states would be appropriate for balancing this trade-off.

D. Comparison with Multivariate HMMs

As mentioned in Section IV-A, we extended our univariate HMM modeling to multivariate HMMs using Gaussian models with maximum force (F_{max}) and contact motion (d) features as the observations. We used scaled features ($S_{F_{max}}$, S_d) for our classification problem. We scaled each feature (f) to a scaled feature (S_f) according to eq. (6) to normalize the values.

$$S_f = (f - mean(f)) / std(f) \tag{6}$$

Fig. 11 shows the results. From the confusion matrix, we observe that the performance has been considerably improved, with an accuracy of 88.57%. Note that the confusion between the fixed and movable categories is much less because

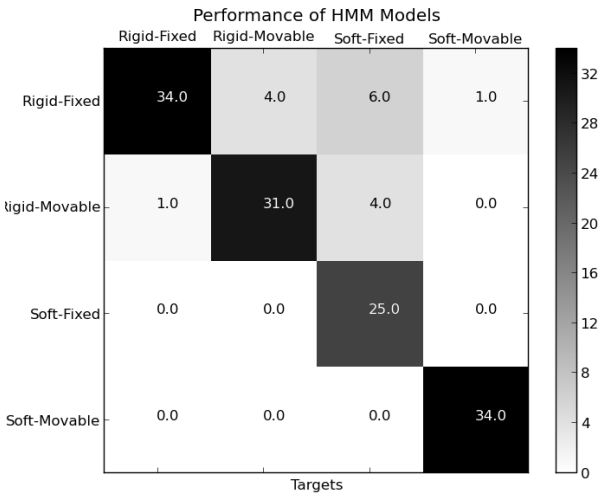

Fig. 11: *Classification into four categories using multivariate HMMs with the robot 'Cody'. In these experiments, the robot moved with stereotyped motion.*

multivariate HMMs can capture the relation between the two features.

E. Comparison with Previous Method

Our previous work in [1] also attempted to categorize objects into these four categories but we used a different method. In our previous work, we extracted the features and converted them to a low-dimensional representation of these feature vectors using PCA. We then used the k-nearest neighbor algorithm (k-NN) to classify a new feature vector based on a set of previously collected feature vectors. To compare our newly developed method with our previous method, developed in [1], we processed the data with our previous algorithm using a dimensionality of three and k=1. This is analogous to HMMs where we selected the most probable state-sequence using the Viterbi algorithm. Fig. 13 shows the results. From the results, we can observe that the previous method performs worse for both one and two features. For a single feature (contact force), the accuracy is 81.43%, whereas for two features (contact force and contact area), the accuracy is 75.71%. As is evident from Fig. 13, multivariate HMMs achieve the highest performance with contact force and contact motion as the observations.

V. SYSTEM MODELING

In Section IV, we presented experimental results using 18 objects made of different materials. These results might not generalize across variations in robot parameters, such as arm compliance or velocity, because of the stereotyped motion of the robot arm during training. In order to gain intuition for this sensing situation, we created an idealized mechanical model of the robot's arm interacting with an object.

We modeled the robot-arm motion similar to our experimental procedure where the robot joint motion is commanded using an equilibrium position of a spring. For our experiments,

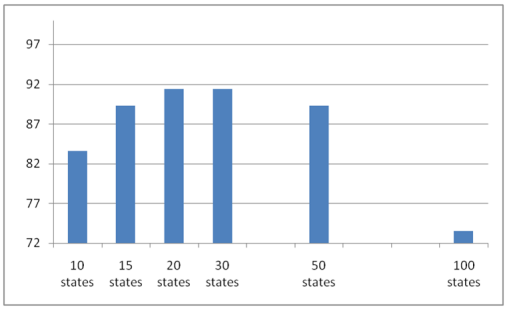

Fig. 12: *As the number of states increases, the accuracy first increases and then decreases. We have shown results up to 100 states with non-uniform intervals, as shown in the figure, to capture the general trend. These results are from experiments in which the robot moved with stereotyped motion.*

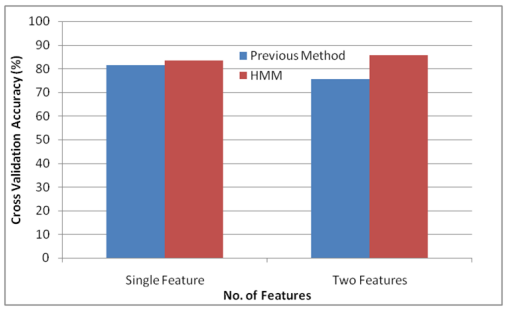

Fig. 13: *Comparison results for experiments with stereotyped motion. The comparison is between HMMs and our previous method in [1]. Multivariate HMMs (Two features) show the best performance.*

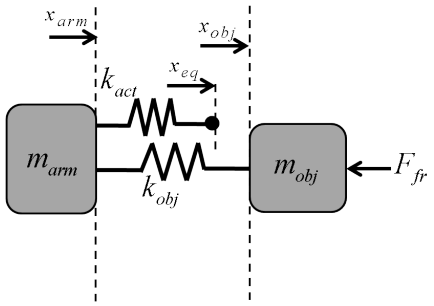

Fig. 14: *Lumped model of our system showing contact between the robot-arm and the object.*

we actuated the robot's elbow joint as described in Section III-B2. This is analogous to a one degree-of-freedom (DOF) motion. Fig. 14 shows a robot-arm of mass m_{arm} making contact with an object of mass m_{obj}. $\boldsymbol{x}_{arm}$ is the position of the robot-arm, $\boldsymbol{x}_{obj}$ is the position of the object, $\boldsymbol{x}_{eq}$ is the equilibrium point of the actuator spring with stiffness k_{act}, and k_{obj} is the object stiffness. F_{fr} is the frictional force acting on the object.

Fig. 15 shows the free-body diagrams of the system in Fig. 14. F_{act} is the force applied on the robot-arm by the actuator, F_{surf} is the force generated when the robot-arm makes contact with the surface of the object, F_{arm} and F_{obj} are the net forces acting on the robot-arm and the object respectively.

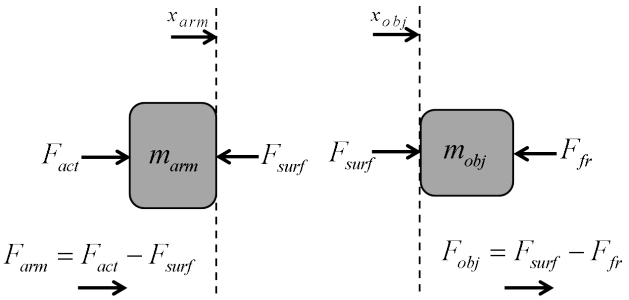

Fig. 15: *Free-body diagrams of the robot-arm and object in contact.*

The force applied by the actuator to the robot-arm, F_{act} is given by eq. (7). The net force on the arm is therefore calculated as in eq. (8). The resultant force on the object is given by eq. (9). Assuming negligible friction at the joint, the resultant position of the robot-arm is calculated as given in eq. (10), where $x_{0_{arm}}$ is the initial position of the arm. The position of object is calculated using eq. (11), where $x_{0_{obj}}$ is the initial position of the object.

$$F_{act} = k_{act}\left(x_{eq} - x_{arm}\right), \tag{7}$$

$$F_{arm} = F_{act} - F_{surf}, \tag{8}$$

$$F_{obj} = F_{surf} - F_{fr}, \tag{9}$$

$$x_{arm} = \iint \left(F_{arm}/m_{arm}\right) dtdt + x_{0_{arm}}, \tag{10}$$

$$x_{obj} = \iint \left(F_{obj}/m_{obj}\right) dtdt + x_{0_{obj}}, \tag{11}$$

The interaction force F_{surf} is calculated using eq. (12), where L_0 is the rest length of the spring.

$$F_{surf} = k_{obj}\left(L_0 - \left(x_{obj} - x_{arm}\right)\right), \tag{12}$$

Please note that the frictional force F_{fr} is calculated differently depending on whether the object is fixed or in motion as shown in eq. (13). If the applied force overcomes the static friction force, the object starts moving. μ_s and μ_k are the coefficients of static and kinetic friction, respectively.

$$F_{fr} = \begin{cases} \mu_s\left(m_{obj}g\right), & fixed \\ \mu_k\left(m_{obj}g\right), & in\ motion \end{cases} \tag{13}$$

Our purpose for creating this model is to better understand the mechanics of robot-object interactions and to get intuition behind the classification problem. As an example, we performed the simulation of a robot-arm making contact with a rigid and a soft object in both fixed and movable conditions. We selected a cube of 10 cm side of uniform density as our object. We selected GRP (Glass Reinforced Plastic) as the material (density = 1500 Kg/m^3 [32]) for a rigid object and soft foam (cushion) as the material (density = 19.22 Kg/m^3 [33]) for a soft object. The resultant mass of the rigid object is

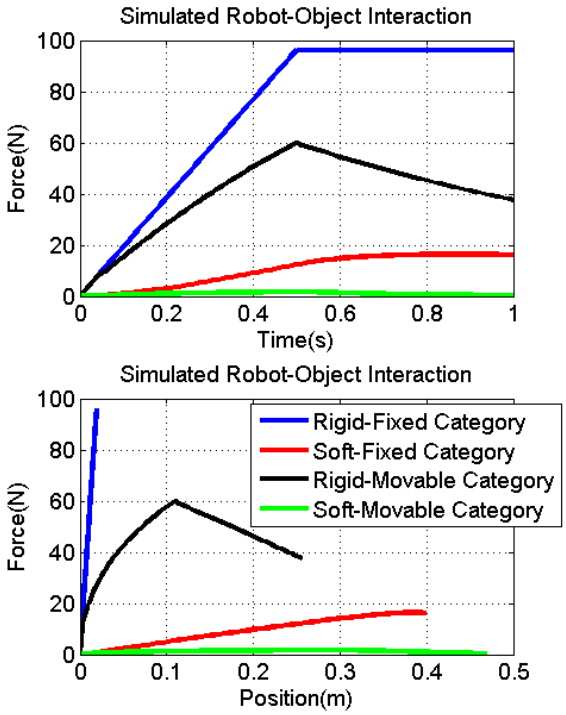

Fig. 16: *Simulation results for force with respect to time and position. The figure shows the forces when the robot is in contact with the object. The simulation runs for 1.0 s. The final goal equilibrium point is at 0.5 m.*

1.5 Kg and the soft object is 20 g. However, for fixed category simulations, we added extra weight to the objects to stop them from moving. The stiffness of the rigid and soft objects are selected as 5000 N/m [32] and 50 N/m [34] respectively. The robot-arm stiffness is selected as 200 N/m based on the elbow joint stiffness of 'Cody' as 25 Nm/rad [2] and a lever arm of roughly 35 cms. The coefficients of static and kinetic friction were selected based on sliding of plastic on plastic for rigid object(0.4 and 0.2 respectively [35]) and foam on plastic for soft object(0.5 and 0.3 respectively [36]). The mass of the 1-DOF robot-arm m_{arm} is 0.5 Kg. Fig. 16 shows the simulation results. Interestingly, the result trends are similar to the actual experimental data in Fig. 5.

VI. Experiments with Varying Robot Parameters

We performed additional experiments to determine whether our algorithm could accommodate variations in robot parameters, such as robot arm compliance and arm velocity. For this new set of experiments, we selected two velocity settings, low=5 deg/s and high=20 deg/s, and two arm stiffness settings, low=2.01 Nm/rad and high=20.1 Nm/rad. We designed this new set of experiments to be similar to the experiments in Section III-B. The robot makes contact with a set of objects incidentally while performing a simple, goal-directed reaching motion. We actuated the robot's elbow joint only and it pushed into soft and rigid objects in fixed and movable conditions with varied arm stiffness (compliance) and different velocities. We performed experiments with eight objects in fixed and movable

Fig. 17: *Set of objects for experiments with variable robot parameters.*

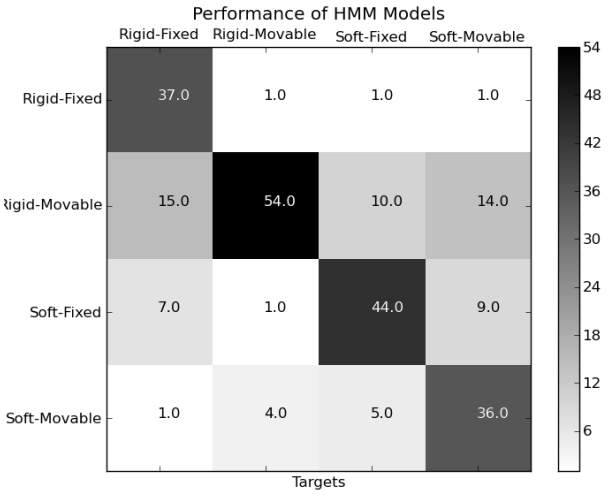

Fig. 18: *The results with multivariate HMMs for experiments with the robot 'Cody'. For these experiments, we varied the robot arm stiffness and velocity.*

conditions (seven in both fixed and movable conditions, one [heavy bucket] in the fixed condition only), as shown in Fig. 17. We repeated the experiments with the two stiffness and velocity settings for four trials each. We collected data for a total of 240 trials [224 (7 objects X 2 stiffness X 2 velocities X 2 conditions X 4 trials) + 16 (1 object X 2 stiffness X 2 velocities X 1 condition X 4 trials)].

A. Additional Experimental Results

We ran this set of 240 trials such that we trained the HMMs with three of the four possible combinations of stiffness and velocity conditions (low-velocity-low-stiffness, low-velocity-high-stiffness, high-velocity-low-stiffness, and high-velocity-high-stiffness) and tested with the other combination to find out how well the results could be extended to different robot parameter conditions. We repeated this procedure for each of the four conditions. Using our previous method (in [1]), the previous algorithm performed poorly, and the accuracy was only 27.5% with a single feature (force) and only 25.42% with two features (force and motion) with the dimensionality of three and k=1. Note that not all trials could be captured up until the time window of 1.2 s because of the varying velocity conditions. In those cases, we extrapolated the data with the mean value for that particular trial to obtain a consistent time window of 1.2 s. With HMMs using a single feature (force), the accuracy improved only slightly to 32.5%, and the performance was still poor. These results are expected because the algorithm does not take into account for the variations in both the robot motion and compliance. However, using multivariate HMMs, the accuracy improved significantly, to 71.25%. Fig. 18 shows the resulting confusion matrix. The majority of the confusion arises between the rigid and soft categories, and we believe that some of this confusion is due to the inconsistency in human labeling because some objects are neither perfectly rigid nor completely compliant. For example, styrofoam is softer than a metal container but we labeled styrofoam as rigid. We labeled hard neoprene rubber as soft though it is much more rigid than a pillow. These results provide evidence that multivariate HMMs can be used to generalize the data-driven inference results to other robots and control methods.

VII. CONCLUSION

In this paper, we developed an object classification algorithm using haptic information obtained from incidental contact of a tactile sensing forearm with objects in the robot's environment. Our algorithm classified objects into four categories: 1) *Rigid-Fixed*, 2) *Rigid-Movable*, 3) *Soft-Fixed*, and 4) *Soft-Movable*. We extracted the temporal trends of maximum force, contact area and contact motion features from the incidental haptic interactions and preprocessed the features to show the information from the onset of contact. We implemented univariate HMMs to model the dynamic interactions with single features and extended the implementation using multivariate HMMs with covarying force and motion features as observations. Our results showed that HMMs are a useful tool to model incidental robot-object interactions and performed better compared to using our previous technique suing PCA + k-NN [1]. We studied the effect of the skin-sensor resolution on the performance of the algorithm. We found that the skin sensor with a higher resolution (384-taxel versus 1-taxel resolution) enhanced the performance of the algorithm when force was used as the observation. We also discussed the effects of the time window of haptic interaction, feature scaling, and the selection of specific features on the overall performance.

In addition, we analyzed the performance of our algorithm with varying robot and object parameter values using a new set of experiments. Multivariate HMMs consistently performed better in all cases with varying robot velocity and compliance parameter values, as is evident from Table III. This finding indicates that our method using multivariate HMMs can continue

TABLE III: Summary of Algorithm Performance.

Task Description	Previous Method [1] with Single Feature	Univariate HMM with Single Feature	Previous Method [1] with Two Features	Multivariate HMM with Two Features
Experiments with Stereotyped Motion	81.43%	83.57%	75.71%	**88.57**%
Experiments with Variable Stiffness and Variable Velocity	27.5%	32.5%	25.42%	**71.25**%

to perform well under circumstances that differ from training, such as different control methods and arm motions.

In total, our results suggest that robots with whole-arm tactile sensing can effectively use data-driven methods to infer useful object properties from incidental contact.

ACKNOWLEDGMENT

We gratefully acknowledge the support from DARPA's Maximum Mobility and Manipulation (M3) program, Contract W911NF-11-1-603, NSF Emerging Frontiers in Research and Innovation (EFRI) 1137229, and NSF Career Award 1150157. We thank Mark Cutkosky and the Stanford Biomimetics and Dexterous Manipulation Lab for their contributions to the forearm tactile skin sensor.

REFERENCES

[1] T. Bhattacharjee, J. M. Rehg, and C. C. Kemp, "Haptic classification and recognition of objects using a tactile sensing forearm," in *IEEE International Conference on Intelligent Robots and Systems (IROS)*, October 2012, pp. 4090–4097.

[2] A. Jain, M. D. Killpack, A. Edsinger, and C. Kemp, "Reaching in clutter with whole-arm tactile sensing," *The International Journal of Robotics Research*, 2013.

[3] T. Bhattacharjee, A. Kapusta, J. M. Rehg, and C. C. Kemp, "Rapid categorization of object properties from incidental contact with a tactile sensing robot arm," in *IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, October 2013.

[4] T. Bhattacharjee, P. M. Grice, A. Kapusta, M. D. Killpack, D. Park, and C. C. Kemp, "A robotic system for reaching in dense clutter that integrates model predictive control, learning, haptic mapping, and planning," in *Proceedings of the 3rd IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) Workshop on Robots in Clutter: Perception and Interaction in Clutter*, 2014.

[5] P. M. Grice, M. D. Killpack, A. Jain, S. Vaish, J. Hawke, and C. C. Kemp, "Whole-arm tactile sensing for beneficial and acceptable contact during robotic assistance," in *13th International Conference on Rehabilitation Robotics (ICORR)*, 2013.

[6] T. Bhattacharjee, A. Jain, S. Vaish, M. D. Killpack, and C. C. Kemp, "Tactile sensing over articulated joints with stretchable sensors," in *IEEE World Haptics Conference (WHC)*, 2013.

[7] L. R. Rabiner, "A tutorial on hidden markov models and selected applications in speech recognition," in *Readings in Speech Recognition*, A. Waibel and K. F. Lee, Eds., Kaufmann, San Mateo, CA, 1990, pp. 267–296.

[8] A. Drimus, G. Kootstra, A. Bilberg, and D. Kragic, "Classification of rigid and deformable objects using a novel tactile sensor," in *in Proceedings of 15th International Conference on Advanced Robotics (ICAR)*, 2011, pp. 427–434.

[9] A. Jain and C. C. Kemp, "Improving robot manipulation with data-driven object-centric models of everyday forces," *Autonomous Robots*, vol. 35, no. 2-3, pp. 143–159, 2013.

[10] V. Chu, I. McMahon, L. Riano, C. G. McDonald, Q. He, J. Perez-Tejada, M. Arrigo, N. Fitter, J. C. Nappo, T. Darrell, and K. J. Kuchenbecker, "Using robotic exploratory procedures to learn the meaning of haptic adjectives," in *Proceedings of International Conference on Robotics and Automation*, 2013.

[11] I. McMahon, V. Chu, L. Riano, G. McDonald, Q. He, J. Perez-Tejada, M. Arrigo, N. Fitter, J. Nappo, T. Darrell, and K. Kuchenbecker, "Robotic learning of haptic adjectives through physical interaction," in *Proceedings of the 2012 Second Workshop on Advances in Tactile Sensing and Touch-based Human-Robot Interaction*, 2012.

[12] V. Sukhoy, R. Sahai, J. Sinapov, and A. Stoytchev, "Vibrotactile recognition of surface textures by a humanoid robot," in *Proceedings of the 2009 Humanoids Workshop: Tactile Sensing in Humanoids Tactile Sensors and Beyond*, 2009.

[13] Y. S. Kim and T. Kesavadas, "Material property recognition by active tapping for fingertip digitizing," in *Proceedings of the Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, March 2006, pp. 133–139.

[14] S. Takamuku, G. Gomez, K. Hosoda, and R. Pfeifer, "Haptic discrimination of material properties by a robotic hand," in *Proceedings of IEEE 6th International Conference on Development and Learning (ICDL)*, 2007, pp. 1–6.

[15] K. Hosoda and T. Iwase, "Robust haptic recognition by anthropomorphic bionic hand through dynamic interaction," in *Proceedings of International Conference on Intelligent Robots and Systems (IROS)*, 2010, pp. 1236–1241.

[16] N. N. A. Charniya and S. V. Dudul, "Sensor for classification of material type and its surface properties using radial basis networks," *IEEE Sensors Journal*, vol. 8, no. 12, pp. 1981–1991, December 2008.

[17] A. Schneider, J. Sturm, and C. Stachniss, "Object identification with tactile sensors using bag-of-features," in *Proceedings of International Conference on Intelligent Robots and Systems (IROS)*, 2009, pp. 243–248.

[18] Z. Pezzementi, E. Plaku, C. Reyda, and G. D. Hager, "Tactile object recognition from appearance information," *IEEE Transactions on Robotics*, vol. 27, no. 3, pp. 473–486, June 2011.

[19] N. Gorges, S. E. Navarro, D. Goger, and H. Worn, "Haptic object recognition using passive joints and haptic key features," in *Proceedings of IEEE International Conference on Robotics and Automation (ICRA)*, May 2010, pp. 2349–2355.

[20] N. Ueda, S. Hirai, and H. T. Tanaka, "Extracting rheological properties of deformable objects with haptic vision," in *Proceedings of the 2004 IEEE International Conference on Robotics and Automation*, April 2004.

[21] B. Frank, R. Schmedding, C. Stachniss, M. Teschner, and W. Burgard, "Learning the elasticity parameters of deformable objects with a manipulation robot," in *Proceedings of International Conference on Intelligent Robots and Systems(IROS)*, 2010, pp. 1877–1883.

[22] ——, "Learning deformable object models for mobile robot navigation using depth cameras and a manipulation robot," in *Proceedings of the Workshop on Advanced Reasoning with Depth Cameras at Robotics, Science and Systems Conference (RSS)*, 2010.

[23] P. K. Allen and K. S. Roberts, "Haptic object recognition using a multi-fingered dexterous hand," in *Proceedings of International Conference on robotics and Automation (ICRA)*, 1989, pp. 342–347.

[24] S. Caselli, C. Magnanini, and F. Zanichelli, "Haptic object recognition with a dexterous hand based on volumetric shape representations," in *Proceedings of the 1994 International Conference on Multisensor Fusion and Integration for Intelligent Systems*, 1994, pp. 280–287.

[25] E. Faldella, B. Fringuelli, D. Passeri, and L. Rosi, "A neural approach to robotic haptic recognition of 3-d objects based on kohonen self-organizing feature map," *IEEE Transactions on Industrial Electronics*, vol. 44, no. 2, pp. 267–269, April 1997.

[26] J. Sinapov, M. Weimer, and A. Stoytchev, "Interactive learning of the acoustic properties of household objects," in *Proceedings of the 2009 IEEE International Conference on Robotics and Automation (ICRA)*, 2009, pp. 2518–2524.

[27] T. Berquist, C. Schenck, U. Ohiri, J. Sinapov, S. Griffith, and A. Stoytchev, "Interactive object recognition using proprioceptive feedback," in *Proceedings of the 2009 IROS Workshop on Semantic Perception for Mobile Manipulation (IROS)*, 2009.

[28] S. Griffith, J. Sinapov, V. Sukhoy, and A. Stoytchev, "A behavior grounded approach to forming object categories: Separating containers from non-containers," *to appear in IEEE Transactions on Autonomous Mental Development*, 2011.

[29] J. Sinapov, T. Bergquist, C. Schenck, U. Ohiri, S. Griffith, and

A. Stoytchev, "Interactive object recognition using proprioceptive and auditory feedback," *International Journal of Robotics Research*, vol. 30, no. 10, pp. 1250–1262, September 2011.
[30] "General Hidden Markov Model Library," http://ghmm.org/.
[31] J. Ulmen, A. Edsinger, and M. Cutkosky, "A highly sensitive, manufacturable, low-cost tactile sensor for responsive robots," in *submitted to IEEE International Conference on Robotics and Automation (ICRA)*, 2012.
[32] "Glass_Fibre_Reinforced_Products," http://www.amiantit.com/media/pdf/brochures/Glass_Fibre_Reinforced_Products/files/Glass_Fibre_Reinforced_Products.pdf.
[33] "Seats & Cushions," http://www.usafoam.com/seat&cushion/seat&cushion.html.
[34] M. K. O'Malley and M. Goldfarb, "The implications of surface stiffness for size identification and perceived surface hardness in haptic interfaces," in *Robotics and Automation, 2002. Proceedings. ICRA'02. IEEE International Conference on*, vol. 2. IEEE, 2002, pp. 1255–1260.
[35] "Coefficient of Friction, Rolling Resistance and Aerodynamics," http://www.tribology-abc.com/abc/cof.htm.
[36] "Coefficients of Friction," http://www.roymech.co.uk/Useful_Tables/Tribology/co_of_frict.htm.

Tapomayukh Bhattacharjee received the B.Tech. degree from the Department of Mechanical Engineering, National Institute of Technology, Calicut, India, and the M.S. degree from the Department of Mechanical Engineering, Korea Advanced Institute of Science and Technology, Daejeon, Korea. He is currently a Robotics Ph.D. Student in the Georgia Institute of Technology, Atlanta. He also worked as a Visiting Scientist with the Interaction and Robotics Research Center, Korea Institute of Science and Technology, Seoul, Korea. His research interests include haptic perception, control systems for robotic manipulation, human-robot interaction, machine learning, and teleoperation systems. He is a member of IEEE, IEEE RAS, and IEEE CSS.

James M. Rehg is a Professor in the School of Interactive Computing at the Georgia Institute of Technology, where he is co-Director of the Computational Perception Lab and is the Associate Director for Research in the Center for Robotics and Intelligent Machines (RIM@GT). He received his Ph.D. from CMU in 1995 and worked at the Cambridge Research Lab of DEC (and then Compaq) from 1995-2001, where he managed the computer vision research group. He received an NSF CAREER award in 2001 and a Raytheon Faculty Fellowship from Georgia Tech in 2005. His research interests include computer vision, medical imaging, robot perception, machine learning, and pattern recognition. Dr. Rehg is currently leading a multi-institution effort to develop the science and technology of Behavior Imaging the capture and analysis of social and communicative behavior using multi-modal sensing, to support the study and treatment of developmental disorders such as autism.

Charles C. Kemp is an Associate Professor at the Georgia Institute of Technology in the Department of Biomedical Engineering. He also has adjunct appointments in the School of Interactive Computing and the School of Electrical and Computer Engineering. He earned a doctorate in Electrical Engineering and Computer Science (2005), an MEng, and BS from MIT. In 2007, he founded the Healthcare Robotics Lab at Georgia Tech (http://healthcare-robotics.com), which focuses on mobile manipulation and human-robot interaction. He is an active member of the Center for Robotics and Intelligent Machines (RIM@GT) and Georgia Tech's multi-disciplinary Robotics Ph.D. program. He has received the 3M Non-tenured Faculty Award, the Georgia Tech Research Corporation Robotics Award, and the NSF CAREER award.